



274.43167X00

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

Applicants: Akira FUJIBAYASHI

Serial No.: 10/784,994

Filed: February 25, 2004

For: METHOD AND APPARATUS FOR MANAGING DIRECT I/O TO  
STORAGE SYSTEMS IN VIRTUALIZATION

**PETITION TO MAKE SPECIAL  
UNDER 37 CFR §1.102(MPEP §708.02)**

**MS Petition**

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

August 4, 2005

Sir:

Applicants hereby petition the Commissioner to make the above-identified application special in accordance with 37 CFR §1.102(d). Pursuant to MPEP §708.02(VIII), Applicants state the following.

**(A) This Petition is accompanied by the fee set forth in 37 CFR §1.17(h).**

The Commissioner is hereby authorized to charge any additional payment due, or to credit any overpayment, to Deposit Account No. 50-1417.

**(B) All claims are directed to a single invention.**

If the Office determines that all claims are not directed to a single invention, Applicant will make an election without traverse as a prerequisite to the grant of special status in conformity with established telephone restriction practice.

**(C) A pre-examination search has been conducted.**

The search was directed towards a network storage virtualization method and a storage as set forth in claims 1-28, particularly claims 1 and 15. According to the present invention a network storage virtualization method in a network storage system and a storage system including first and second network storage devices and a virtualization apparatus is provided. The present invention includes allowing a client connected via a network to access the network storage devices as one virtual network storage system; and permitting said client to access the network storage devices separate from the virtualized network storage system. Under additional aspects, the present invention provides that the first network storage device includes a disk interface (I/F) which is connected to an apparatus having a function of forming the network storage devices as one virtual network storage device, and that the second network storage device includes a disk I/F which is connected to the apparatus having a function of forming the network storage devices as one virtual network storage device and a host I/F which is connectable to an external device so as to permit the external device to access the network storage devices separate from the apparatus having a function of forming the network storage devices as one virtual network storage device. The external device is permitted to access a secondary volume which is a copy of a primary volume and not permitted to access volumes other than the secondary volume.

The search of the above features was conducted in the following areas:

<u>Class</u>	<u>Subclasses</u>
707	200, 204
709	212-219
711	112, 114, 141, 161, 162
714	4, 6, 112, 114, 147, 148, 170, 173, 202, 203

Additionally, a computer database search was conducted on the USPTO systems EAST and WEST.

**(D) The following is a list of the references deemed most closely related to the subject matter encompassed by the claims:**

<u>U.S. Patent Application Publication No.</u>	<u>Inventor(s)</u>
2002/0004857	Arakawa et al
2002/0112113	Karpoff et al
2003/0055972	Fuller et al
2003/0115218	Bobbitt et al
2003/0140051	Fujiwara et al
2003/0191810	Muhlestein et al
2003/0220985	Kawamoto et al
<u>Foreign Documents</u>	<u>Inventor</u>
CA 2344597	Arndt
EP 1178407	Reuter et al.

A copy of each of these references (as well as other references uncovered during the search) is enclosed in an accompanying IDS.

**(E) It is submitted that the present invention is patentable over the references for the following reasons.**

It is submitted that the cited references, whether taken individually or in combination with each other, fail to teach or suggest the invention as claimed. In

particular, the cited references, at a minimum, fail to teach or suggest as recited in the claims:

a first feature of the present invention as recited in each of independent claims 1 and 15 of permitting the client to access the network storage devices separate from the virtualized network storage system.

Further, the cited references fail to teach or suggest the above noted features of the present invention when taken in combination with the other limitations recited in the claims.

The references considered most closely related to the claimed invention are briefly discussed below:

Arndt (CA No. 2344597) discloses a logically partitioned data processing system in which shared resources are emulated to provide each partition a separate copy of the shared resource. In one embodiment, the logically partitioned data processing system includes a plurality of logical partitions, a plurality of operating systems executing within the data processing system and a plurality of assignable resources. Each of the plurality of operating systems is assigned to a separate one of the plurality of logical partitions, such that no more than one operating system is assigned to any given logical partition. Each of the plurality of assignable resources is assigned to a single one of the plurality of logical partitions. The logically partitioned data processing system also includes a hypervisor. The hypervisor emulates shared resources, such as an operator panel and a system console, and provides a virtual copy of these shared resources to each of the plurality of logical partitions. Thus, each operating

system executing within data processing system 300 may access only those I/O units that are within its logical partition. (See, e.g., Abstract and p. 6, lines 8-20).

However, unlike the present invention, Arndt does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Arndt at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Arakawa (U.S. Patent Application Publication No. 20020004857) discloses techniques, including a method and system, for relocating data between storage systems. In one embodiment of the present invention a host collects usage information from a plurality of storage systems, and determines the relocation destination LU for data stored in the LU to be relocated. The host alters an LU logical position name table that determines matching between the logical position names of data and LUs. It also carries out data relocation between storage subsystems by shifting data stored in an origin LU to be relocated to a destination LU. In another embodiment relocation of files is provided. Disclosed is a network storage system for providing a plurality of distributed storage

systems viewed by a client as one virtual storage system. (See, e.g., Abstract and claim 33).

However, unlike the present invention, Arakawa does not teach or suggest a network storage virtualization method Arndt does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Arakawa at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Reuter (European Patent No. 1178407) discloses a system and method for creating virtualized storage in a storage area network using distributed table-driven input/output mapping. The present invention distributes the virtualization mapping in multiple parallel, mapping agents that are separate from a controller. This allows the performance-sensitive mapping process to be parallelized and distributed optimally for performance, while the control of the mapping may be located in a controller chosen for optimal cost, management, and other implementation practicalities. The mapping agents store the virtual mapping tables in volatile memory, substantially reducing the cost and complexity of

implementing the mapping agents. The controller is responsible for persistent storage of mapping tables, thereby consolidating the costs and management for persistent mapping table storage in a single component. Distributed virtualization also allows the controller to manage multiple virtual disks used by multiple host systems, and allows a single virtual disk to be shared by multiple host systems. The mapping agents preferably do not interact only with other mapping agents, thereby improving the scalability of the virtual storage system and the virtual storage system's tolerance of component failures. (See, e.g., Abstract and paragraphs 12-14).

However, unlike the present invention, Reuter does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Reuter at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Karpoff (U.S. Patent Application Publication No. 20020112113) discloses storage virtualization systems and methods that allow customers to manage storage as a utility rather than as islands of storage which are independent of

each other. A demand mapped virtual disk image of up to an arbitrarily large size is presented to a host system. The virtualization system allocates physical storage from a storage pool dynamically in response to host I/O requests, e.g., SCSI I/O requests, allowing for the amortization of storage resources through a disk subsystem while maintaining coherency amongst I/O RAID traffic. In one embodiment, the virtualization functionality is implemented in a controller device, such as a controller card residing in a switch device or other network device, coupled to a storage system on a storage area network (SAN). The resulting virtual disk image that is observed by the host computer is larger than the amount of physical storage actually consumed. Remapping the storage allows the back-end storage to be managed without consumer impact and multiple back-end partitions to be combined to provide a single virtual image. The disk image of the present invention presents potentially a very large image to the consumer to isolate him from volume resizing issues and to allow easy consumption. This image may be supported by a management system that provides the ability to control consumption and growth rates as well as maintain core system processes such as creating, deleting and mounting other candidate disks. (See, e.g., Abstract and paragraph 16).

However, unlike the present invention, Karpoff does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage



devices separate from the virtualized network storage system as in the present invention.

More particularly, Karpoff at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Fuller (U.S. Patent Application Publication No. 20030055972) discloses methods and systems for shared storage virtualization. In a shared storage infrastructure accessible by more than one customer, the shared storage infrastructure is separated into a plurality of logical storage areas. A customer is associated with at least one logical storage area within the shared storage infrastructure. The customer and its logical storage areas are associated with at least one port in the shared storage infrastructure. Each customer has access only to the logical storage areas associated with the customer and cannot access the logical storage areas of any other customer. The shared storage infrastructure provides policy based storage management to each customer, that is, each customer appears to have access to a non-shared storage subsystem. By providing a virtualization layer, a customer may be given the ability to manage its own storage needs. At step 410, storage area 360 is partitioned into a plurality of smaller storage areas or logical volume units and at least one of the logical volume units is associated with customer 310. At step 420, the logical volume associated with customer 310 is associated with the at least one unique

port connection associated with customer 310 in step 400. At step 430, router 350 receives a request to access the unique port connection for customer 310 from customer 310. (See, e.g., Abstract and paragraphs 12 and 55).

However, unlike the present invention, Fuller does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Fuller at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Bobbitt (U.S. Patent No. 2003/0115218) discloses a virtual file system and method. The system architecture enables a plurality of underlying file systems running on various file servers to be "virtualized" into one or more "virtual volumes" that appear as a local file system to clients that access the virtual volumes. The system also enables the storage spaces of the underlying file systems to be aggregated into a single virtual storage space, which can be dynamically scaled by adding or removing file servers without taking any of the file systems offline and in a manner transparent to the clients. This functionality is enabled through a software "virtualization" filter on the client that intercepts file

system requests and a virtual file system driver on each file server. The system also provides for load balancing file accesses by distributing files across the various file servers in the system, through migration of data files between servers. The invention is a method for virtualizing a plurality of underlying file systems hosted on one or more file servers, wherein each file system includes a storage space in which a plurality of data files may be stored comprising: aggregating the storage spaces of said plurality of file systems into a single virtual storage space that is exported as a virtual volume that appears to clients as local file system. (See, e.g., Abstract and claim 1).

However, unlike the present invention, Bobbitt does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Bobbitt at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Fujiwara (U.S. Patent Application Publication No. 20030140051) discloses a method of virtualizing a plurality of network storages into a single-view file system for a client and obtains information about an object's storage location

without inquiring of a resource manager. The method determines a network storage that should store an object based on a hash value generated from a logical identifier for the object to be accessed. While accessing the object, the method computes the hash value for the object's logical identifier to determine a network storage that should process an access request. Consequently, the method eliminates the need to inquire the object's storage location when the access request is issued. Further, the method eliminates the need to maintain the object's storage location when the number of servers is increased or decreased, thus providing easy management. The invention provides a virtualized single-view method for a plurality of network storages so as to be capable of uniquely determining a network storage for storing an object to be accessed from a logical identifier for the object to be accessed. Another object of the present invention is to provide a virtualized single-view method for a plurality of network storages with improved network storage expandability. (See, e.g., Abstract and paragraphs 9-11).

However, unlike the present invention, Fujiwara does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Fujiwara at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of

independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Muhlestein (U.S. Patent Application Publication No. 20030191810) discloses an architecture that provides the ability to create and maintain multiple instances of virtual servers, such as virtual filers (vfilers), within a server, such as a filer. A vfiler is a logical partitioning of network and storage resources of the filer platform to establish an instance of a multi-protocol server. Each vfiler is allocated a subset of dedicated units of storage resources, such as volumes or logical sub-volumes (qtrees), and one or more network address resources. Each vfiler is also allowed shared access to a file system resource of a storage operating system. To ensure controlled access to the allocated and shared resources, each vfiler is further assigned its own security domain for each access protocol. A vfiler boundary check is performed by the file system to verify that a current vfiler is allowed to access certain storage resources for a requested file stored on the filer platform. When the physical interfaces 218 and their associated links 180 are aggregated as a single virtual interface 190, all of the physical interfaces respond to only one MAC address. That is, the physical interfaces 218 are organized into one virtual "pipe" having one logical interface that is assigned a common MAC address. (See. e.g., Abstract and paragraph 43).

However, unlike the present invention, Muhlestein does not teach or suggest a network storage virtualization method or storage system which allows

a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Muhlestein at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Kawamoto (U.S. Patent Application Publication No. 20030220985) discloses a virtualizing file system view method for virtualizing one or more network storage devices into a virtualized file system view network storage system wherein destination network storage information of stored files is compactly held regardless of the number of files, and files are separated into one or more file groups and the file groups managed in destination network storage units. Until now, managing network storage unit information in individual files was necessary however the virtualizing file system view method reduces the management information that must be held and efficiently utilizes network storage capacity without holding destination network storage information in individual files. The cost of rewriting information is also lowered during structural changes such as adding or deleting network storage units since storage destination network information can be rewritten in file groups. The present invention has the object of providing a method for virtualizing multiple network

storage units into a single network storage system for making management information for the file storage destination more compact and that does not place a strain on the data storage area. Another object of the present invention is to provide a method for virtualizing multiple network storage units into a single network storage system that reduces the overhead of the task of rewriting management information during file transfer caused by structural changes in network storage, and that permits fast and flexible structural changes. (See, e.g., Abstract and paragraphs 14-16).

However, unlike the present invention, Kawamoto does not teach or suggest a network storage virtualization method or storage system which allows a client connected via a network to access the network storage devices as one virtual network storage system, wherein the client can access the network storage devices separate from the virtualized network storage system as in the present invention.

More particularly, Kawamoto at a minimum does not teach or suggest the above described first feature of the present invention as recited in each of independent claims 1 and 15, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Therefore, since the cited references at a minimum fail to teach or the above described first feature of the present invention as recited in each of independent claim 1 and 15, and further fail to teach or suggest these features in combination with the other limitations recited in each of the independent claims, it

is submitted that all of the claims are patentable over the cited references whether said references are taken individually or in combination with each other.

**(F) Conclusion**

Applicant has conducted what it believes to be a reasonable search, but makes no representation that "better" or more relevant prior art does not exist. The United States Patent and Trademark Office is urged to conduct its own complete search of the prior art, and to thoroughly examine this application in view of the prior art cited herein and any other prior art that the United States Patent and Trademark Office may locate in its own independent search. Further, while Applicant has identified in good faith certain portions of each of the references listed herein in order to provide the requisite detailed discussion of how the claimed subject matter is patentable over the references, the United States Patent and Trademark Office should not limit its review to the identified portions but rather, is urged to review and consider the entirety of each reference, and not to rely solely on the identified portions when examining this application.

In view of the foregoing, Applicant requests that this Petition to Make Special be granted and that the application undergo the accelerated examination procedure set forth in MPEP 708.02 VIII.



**(G) Fee (37 C.F.R. 1.17(h))**

The fee required by 37 C.F.R. § 1.17(h) is to be paid by:

☒ the Credit Card Payment Form (attached) for \$130.00.


☐ charging Account \_\_\_\_\_ the sum of \$130.00.

A duplicate of this petition is attached.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.,  
Deposit Account No. 50-1417 (274.43167X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.

  
\_\_\_\_\_  
Carl I. Brundidge  
Reg. No. 29,621

CIB/jdc  
(703) 684-1120

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

**PETITION FEE**

Under 37 CFR 1.17(f), (g) &amp; (h)

**TRANSMITTAL**

(Fees are subject to annual revision)

Send completed form to: Commissioner for Patents  
P.O. Box 1450, Alexandria, VA 22313-1450

Application Number	10/784,994
Filing Date	February 25, 2004
First Named Inventor	Akira FUJIBAYASHI
Art Unit	
Examiner Name	
Attorney Docket Number	274.43167X00

Enclosed is a petition filed under 37 CFR §1.102(d) that requires a processing fee (37 CFR 1.17(f), (g), or (h)). Payment of \$ 130.00 is enclosed.

This form should be included with the above-mentioned petition and faxed or mailed to the Office using the appropriate Mail Stop (e.g., Mail Stop Petition), if applicable. For transmittal of processing fees under 37 CFR 1.17(i), see form PTO/SB/17i.

**Payment of Fees** (small entity amounts are NOT available for the petition (fees))

- ☒ The Commissioner is hereby authorized to charge the following fees to Deposit Account No. 50-1417:
- ☐ petition fee under 37 CFR 1.17(f), (g) or (h) ☒ any deficiency of fees and credit of any overpayments
- Enclose a duplicative copy of this form for fee processing.

☐ Check in the amount of \$ \_\_\_\_\_ is enclosed.☒ Payment by credit card (From PTO-2038 or equivalent enclosed). Do not provide credit card information on this form.**Petition Fees under 37 CFR 1.17(f): Fee \$400 Fee Code 1462**

For petitions filed under:

- § 1.53(e) - to accord a filing date.
- § 1.57(a) - to according a filing date.
- § 1.182 - for decision on a question not specifically provided for.
- § 1.183 - to suspend the rules.
- § 1.378(e) for reconsideration of decision on petition refusing to accept delayed payment of maintenance fee in an expired patent.
- § 1.741(b) - to accord a filing date to an application under §1.740 for extension of a patent term.

**Petition Fees under 37 CFR 1.17(g): Fee \$200 Fee code 1463**


For petitions filed under:

- §1.12 - for access to an assignment record.
- §1.14 - for access to an application.
- §1.47 - for filing by other than all the inventors or a person not the inventor.
- §1.59 - for expungement of information.
- §1.103(a) - to suspend action in an application.
- §1.136(b) - for review of a request for extension of time when the provisions of section 1.136(a) are not available.
- §1.295 - for review of refusal to publish a statutory invention registration.
- §1.296 - to withdraw a request for publication of a statutory invention registration filed on or after the date the notice of intent to publish issued.
- §1.377 - for review of decision refusing to accept and record payment of a maintenance fee filed prior to expiration of a patent.
- §1.550(c) - for patent owner requests for extension of time in ex parte reexamination proceedings.
- §1.956 - for patent owner requests for extension of time in inter partes reexamination proceedings.
- § 5.12 - for expedited handling of a foreign filing license.
- § 5.15 - for changing the scope of a license.
- § 5.25 - for retroactive license.

**Petition Fees under 37 CFR 1.17(h): Fee \$130 Fee Code 1464**

For petitions filed under:

- §1.19(g) - to request documents in a form other than that provided in this part.
- §1.84 - for accepting color drawings or photographs.
- §1.91 - for entry of a model or exhibit.
- §1.102(d) - to make an application special.
- §1.138(c) - to expressly abandon an application to avoid publication.
- §1.313 - to withdraw an application from issue.
- §1.314 - to defer issuance of a patent.

Name (Print/Type)	Carl L. Brundidge	Registration No. (Attorney/Agent)	29,621
Signature		Date	August 4, 2005

This collection of information is required by 37 CFR 1.114. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.